

Graphical Methods for Inequality Constraints in Marginalized DAGs

Robin J. Evans

September 14, 2012

Abstract

We present a graphical approach to deriving inequality constraints for directed acyclic graph (DAG) models, where some variables are unobserved. In particular we show that the observed distribution of a discrete model is always restricted if any two observed variables are neither adjacent in the graph, nor share a latent parent; this generalizes the well known instrumental inequality. The method also provides inequalities on interventional distributions, which can be used to bound causal effects. All these constraints are characterized in terms of a new graphical separation criterion, providing an easy and intuitive method for their derivation.

1 Introduction

Models based on directed acyclic graphs (DAGs) are commonly used for causal inference on account of their simple to understand conditional independence constraints, and the intuitive appeal of using arrows to display causal dependences. If all the variables in a DAG are observed then causal quantities of interest are typically point identified, and derivable in terms of conditional probabilities. However, it is common for some variables to be unobservable, possibly representing confounding factors which may bias inference; in this case we can only observe the marginal distribution over the remaining variables.

The models which result from the marginalization of a DAG are much less well understood and, unlike DAGs, are not described merely in terms of conditional independence constraints. In particular, causal effects may not be point identified, and we can only hope for inequality constraints describing the range of possible values.

Existing methods for deriving bounds on observed distributions are either specific to a particular model Pearl (1995); Balke and Pearl (1997), or computationally intensive and lacking the intuitiveness of a graphical approach Bonet (2001); Kang and Tian (2006). See Ramsahai (2012) for an approach which is graphical in spirit, but uses computationally difficult variable elimination methods. In this paper we take steps to remedy these problems by providing a simple

graphical separation criterion for determining the existence of constraints, and for constructing them explicitly.

The remainder of the paper is organised as follows: §2 introduces DAGs and related terminology and notation. §3 gives a new method for deriving known constraints on the observed distribution of the instrumental variables model, and related causal effects. §4 applies these methods to give new constraints for general DAG models, and §5 contains examples. A discussion is found in §6, and longer proofs are in an appendix.

2 Graphical Models

A *directed graph* \mathcal{G} is a set of vertices \mathbf{V} , with a collection of ordered pairs of distinct vertices, or *edges*, \mathcal{E} . If $(X, Y) \in \mathcal{E}$ we write $X \rightarrow Y$, and say that X is a *parent* of Y . The set of parents of Y is denoted $\text{pa}_{\mathcal{G}}(Y)$. A *path* is a sequence of adjacent edges in a graph, without repetition of vertices; for example, the graph in Figure 1(a) contains the path $\pi_1 : Z \rightarrow X \leftarrow U \rightarrow Y$. A path is *directed* from X to Y if all the arrows point away from X and towards Y . If there is a directed path from X to Y we say that Y is a *descendant* of X , and X an *ancestor* of Y . A directed graph is *acyclic* if there is no directed path from a vertex to itself; such an object is called a directed acyclic graph (DAG).

We associate each vertex X with a random variable under some multivariate distribution P ; let P admit a density f . For convenience, in what follows we will use X to denote both the vertex and the random variable, and similarly use operators and bold face letters (e.g. $\text{pa}_{\mathcal{G}}(X)$, \mathbf{C}) to refer to both a set of vertices and the associated vector of random variables. The *factorization criterion* for DAGs says that P is in the model corresponding to the DAG \mathcal{G} if the joint density factorizes as $\prod_{V \in \mathbf{V}} f(V \mid \text{pa}_{\mathcal{G}}(V))$.

Internal vertices on a path with two adjacent arrowheads are called *colliders* on the path; other internal vertices are non-colliders. On the path π_1 , X is a collider, and U a non-collider. A path π from X to Y is *blocked* given a set of vertices \mathbf{C} if there is a non-collider on π in \mathbf{C} , or a collider on π which is not an ancestor of any vertex in \mathbf{C} .

We say that two sets of vertices \mathbf{A} and \mathbf{B} are *d-separated* given a set of vertices \mathbf{C} , if every path from any vertex in \mathbf{A} to any vertex in \mathbf{B} is blocked by \mathbf{C} . A probability distribution P obeys the *global Markov property* for a DAG \mathcal{G} if whenever \mathbf{A} and \mathbf{B} are d-separated by \mathbf{C} in \mathcal{G} , then $\mathbf{A} \perp\!\!\!\perp \mathbf{B} \mid \mathbf{C} [P]$.

It is well known that d-separation is equivalent to the factorization criterion Verma and Pearl (1988). In particular, all constraints implied by a DAG on fully observed random variables can be interpreted as conditional independences.

Assigning a causal interpretation to a DAG model requires extra assumptions, in particular that the system under observation is stable under interventions with respect to the graph. We will denote an intervention to fix $X = x$ by $\text{do}(X = x)$, or $\text{do}(x)$ for short; graphically this may be represented by removing the edges of the form $Y \rightarrow X$, so that X has no parents in the new graph. The density $f(\mathbf{V} \mid \text{do}(x))$ is given by dividing the joint density f by $f(x \mid \text{pa}_{\mathcal{G}}(X))$




Figure 1: (a) The instrumental variables (IV) model; U is unobserved. (b) The IV model with the effect of X on Y removed.

and multiplying by the indicator function $\mathbb{1}_{\{X=x\}}$. See Pearl (2009) for details.

If some of the variables in a DAG are unobserved, we may be interested in the implications of the underlying graph for the observable margin. Let $\mathbf{U} \subset \mathbf{V}$ denote the set of latent or unobservable vertices; the observable margin is then

$$\int_{\mathbf{U} \in \mathbf{U}} \prod_{V \in \mathbf{V}} f(V \mid \text{pa}_{\mathcal{G}}(V)) d\mathbf{U}. \quad (1)$$

The marginal distribution over the observed variables is completely identifiable, but some of the structure of the underlying graph may be impossible to determine in the presence of latent variables. We will make no assumption about the state space of the latent variables, since these are unobserved. Some conditional independences may still be observable, but other kinds of constraint also arise, including *Verma constraints* Verma and Pearl (1991), and inequalities on the observed distribution (see next section).

Without loss of generality we will assume that none of the latent variables have any parents.

3 Instrumental Variables

Perhaps the most thoroughly studied causal DAG model is the *instrumental variables* model, represented in Figure 1(a). It arises naturally in randomized trials with imperfect compliance, in which Z represents a randomized treatment assignment, X the treatment actually taken by the subject, and Y an outcome; U represents unmeasured confounding factors which may affect both the probability of the subject taking the treatment and the outcome of interest, so that naïve estimators of the effect of X on Y will be biased.

The graph encodes (amongst other assumptions) that the assignment Z does not affect the outcome Y other than through the treatment X . This is known as the exclusion restriction, and is important for assessing the effect of X on Y ; implications of the exclusion restriction which can be subjected to an empirical test are therefore very useful.

Making no assumptions about the character of U , and if X is continuous, the observable margin is unconstrained Bonet (2001). However, if the observed

variables have finite and discrete state spaces, then the observed distribution obeys the *instrumental inequality* of Pearl (1995):

$$\max_x \sum_y \max_z p(x, y | z) \leq 1; \quad (2)$$

here $p(x, y | z)$ is used to denote $P(X = x, Y = y | Z = z)$. This restriction can be used to falsify the IV model. Pearl's proof of the inequality is model specific, and it is not clear how it might be applied to other graphs. Below we present a new approach to the derivation of (2), and a more graphical interpretation of its meaning; as we shall see, this method can be adapted to many other DAG models, and provides some causal constraints.

Proposition 3.1. *Let P be a probability distribution over three random variables Z , X and Y , taking values in discrete sets \mathcal{Z} , \mathcal{X} and \mathcal{Y} respectively. Then P obeys the IV model only if for each $\xi \in \mathcal{X}$, the collection of conditional probabilities $(p(\xi, y | z), y \in \mathcal{Y}, z \in \mathcal{Z})$ is compatible with a distribution under which $Y \perp\!\!\!\perp Z$.*

In other words, only if for each $\xi \in \mathcal{X}$ there exists a distribution P^ such that $Y \perp\!\!\!\perp Z [P^*]$, and $p^*(\xi, y | z) = p(\xi, y | z)$ for each $y \in \mathcal{Y}$ and $z \in \mathcal{Z}$.*

This condition implies the instrumental inequality (2).

Proof. Suppose that P is in the IV model. Then

$$p(x, y | z) = \int f(u) p(x | u, z) p(y | u, x) du;$$

construct a distribution P^* by

$$p^*(x, y | z) = \int f(u) p(x | u, z) p(y | u, \xi) du.$$

Under P^* , the effect of X on Y has been broken, because Y behaves as though $X = \xi$ regardless of its actual value. P^* obeys the factorization criterion with respect to the graph in Figure 1(b); thus $Y \perp\!\!\!\perp Z [P^*]$, and by construction $p^*(\xi, y | z) = p(\xi, y | z)$ for each $y \in \mathcal{Y}$ and $z \in \mathcal{Z}$.

To see that this implies (2), first note that the independence is equivalent to

$$\begin{aligned} p^*(y | z) &= p^*(y | z') \\ p(\xi, y | z) + \sum_{x \neq \xi} p^*(x, y | z) &= p(\xi, y | z') + \sum_{x \neq \xi} p^*(x, y | z') \end{aligned}$$

for each $y \in \mathcal{Y}$, $z, z' \in \mathcal{Z}$. Suppose we are given the probabilities $p(\xi, y | z)$ and asked to construct a distribution P^* satisfying these equations. Since all the quantities are positive, and this equality holds for each z' , we have

$$\max_{z'} p(\xi, y | z') - p(\xi, y | z) \leq \sum_{x \neq \xi} p^*(x, y | z).$$




Figure 2: The model from Figure 1(a) with a possible effect of Z on Y added.

However the sum of the quantities on the RHS over y cannot be greater than $1 - p(\xi | z) = 1 - \sum_y p(\xi, y | z)$, so

$$\sum_y \left(\max_{z'} p(\xi, y | z') - p(\xi, y | z) \right) \leq 1 - \sum_y p(\xi, y | z)$$

$$\sum_y \max_{z'} p(\xi, y | z') \leq 1.$$

Applying this to each ξ gives (2). \square

Remark 3.2. *Whilst these inequalities are not new, the importance of the above result lies in the proof technique; we will see in the next section that it generalizes to many other DAG models, giving novel results.*

The instrumental inequality is exact when X , Y and Z are binary, but insufficient if Z takes three states Bonet (2001). The sufficient bounds are difficult to derive without using computationally intensive linear programming techniques and Fourier-Motzkin elimination, which become infeasible for moderately sized state spaces.

3.1 Causal bounds on the IV model

We next try to invert the problem and ask *how much* effect Z can have on Y given the observed distribution. In some sense we are trying to quantify the strength of the dashed arrow in Figure 2. A suitable measure is the *average controlled direct effect* (ACDE) of Z on Y , controlling for $X = x$; this is defined for binary Z , X and Y as

$$\text{ACDE}_{Z \rightarrow Y}(x) \equiv p(y_1 | \text{do}(z_1, x)) - p(y_1 | \text{do}(z_0, x)).$$

Here y_1 is a shorthand for $\{Y = 1\}$, whilst x means $\{X = x\}$, etc. Generalizations to non-binary state spaces are also possible Cai et al. (2008). Note that $\text{ACDE}_{Z \rightarrow Y}(x) = 0$ for each x if $Z \not\rightarrow Y$. For the DAG in Figure 2,

$$p(y | \text{do}(z, x)) = \int_u f(u) p(y | x, z, u) du,$$

which is not identified. However, constructing P^* as above,

$$\begin{aligned}
p(y \mid \text{do}(z, \xi)) &= \int_u f(u) p^*(y \mid z, u) du \\
&= p^*(y \mid z) \\
&= p(y, \xi \mid z) + \sum_{x \neq \xi} p^*(y, x \mid z) \\
&\leq p(y, \xi \mid z) + 1 - p(\xi \mid z).
\end{aligned}$$

Also $p(y \mid \text{do}(z, \xi)) \geq p(y, \xi \mid z)$, so

$$\begin{aligned}
\text{ACDE}_{Z \rightarrow Y}(x) &\leq p(y_1, x \mid z_1) + 1 - p(x \mid z_1) - p(y_1, x \mid z_0) \\
&= 1 - p(y_0, x \mid z_1) - p(y_1, x \mid z_0),
\end{aligned}$$

and similarly

$$\text{ACDE}_{Z \rightarrow Y}(x) \geq p(y_1, x \mid z_1) + p(y_0, x \mid z_0) - 1.$$

Note that the ACDE bounds include zero if and only if the instrumental inequality (2) is satisfied. These bounds were derived by Cai et al. (2008) using linear programming, and shown to be tight. In the next section we will extend this method to other graphs.

4 Other Models

Just as d-separation provides a graphical criterion for finding observable conditional independences, we now provide a graphical criterion for finding observable inequality constraints. For a DAG \mathcal{G} with vertex set \mathbf{V} and edge set \mathcal{E} , define the *induced subgraph* $\mathcal{G}_{\mathbf{W}}$ for $\mathbf{W} \subset \mathbf{V}$ as the DAG with vertex set \mathbf{W} and edge set $\mathcal{E} \cap (\mathbf{W} \times \mathbf{W})$.

Now we define our new separation criterion: let \mathbf{A} , \mathbf{B} , \mathbf{C} and \mathbf{D} be disjoint sets of observed vertices. \mathbf{A} and \mathbf{B} are *e-separated* (extended d-separation) given \mathbf{C} after deletion of \mathbf{D} in \mathcal{G} , if \mathbf{A} and \mathbf{B} are d-separated by \mathbf{C} in $\mathcal{G}_{\mathbf{V} \setminus \mathbf{D}}$. In other words, if we remove the vertices in \mathbf{D} from the graph, then \mathbf{A} and \mathbf{B} are d-separated by \mathbf{C} .

For example, in the graph in Figure 1(a), Z and Y are e-separated after deletion of X . The following lemma gives an alternative characterization of e-separation which will prove useful. Its proof is elementary, and omitted for brevity.

Lemma 4.1. *Let \mathcal{G} be a DAG, and let \mathcal{G}^* be the DAG formed from \mathcal{G} by removing all edges which are oriented away from some vertex in \mathbf{D} (i.e. of the form $D \rightarrow E$ for $D \in \mathbf{D}$). Then \mathbf{A} is e-separated from \mathbf{B} by \mathbf{C} after deletion of \mathbf{D} in \mathcal{G} if and only if \mathbf{A} is d-separated from \mathbf{B} by \mathbf{C} in \mathcal{G}^* .*

Graphs formed by removing the edges emanating from vertices form a part of Pearl’s do-calculus Pearl (2009). The node-splitting method in Robins et al. (2006) is also related.

Suppose now that we are interested in the detecting the presence or absence of the edge $X \rightarrow Y$ in a general graph, and in estimating the strength of the (direct) causal effect of X on Y . We first show that if X and Y are not directly confounded with each other, which is to say that they do not share a latent parent, then falsifiable constraints (such as the instrumental inequality) for the absence of the edge $X \nrightarrow Y$ always exist.

Theorem 4.2. *Let \mathcal{G} be a DAG, and let \mathbf{A} , \mathbf{B} , \mathbf{C} and \mathbf{D} be disjoint sets of observable vertices such that no vertex in \mathbf{C} is a descendant of any in \mathbf{D} . If \mathbf{A} is e-separated from \mathbf{B} by \mathbf{C} after deletion of \mathbf{D} , then for any fixed value $\mathbf{D} = \mathbf{d}$, the conditional probabilities $p(\mathbf{a}, \mathbf{b}, \mathbf{d} | \mathbf{c})$ must be compatible with a distribution P^* in which $\mathbf{A} \perp\!\!\!\perp \mathbf{B} | \mathbf{C} [P^*]$.*

If in addition no vertex in \mathbf{A} is a descendant of any element of \mathbf{D} , then the probabilities $p(\mathbf{b}, \mathbf{d} | \mathbf{a}, \mathbf{c})$ must be compatible with a distribution P^ in which $\mathbf{A} \perp\!\!\!\perp \mathbf{B} | \mathbf{C} [P^*]$.*

Corollary 4.3. *Let \mathcal{G} be a DAG containing observable vertices X, Y , which do not share a latent parent nor are joined by an edge; let \mathcal{G}' be equal to \mathcal{G} , except that $X \rightarrow Y$ in \mathcal{G}' . Then if the observed variables in the graphs are discrete, the model defined by the observed margin of \mathcal{G}' is strictly larger than the one defined by \mathcal{G} .*

Proof. Under the conditions given, we can apply Theorem 4.2 to \mathcal{G} with $\mathbf{A} = \{X\}$, $\mathbf{B} = \{Y\}$, $\mathbf{C} = \emptyset$ and $\mathbf{D} = \mathbf{V} \setminus (\mathbf{U} \cup \{X, Y\})$.

To see that this implies a constraint, consider a distribution in which all vertices other than X and Y are completely independent, and $P(\mathbf{D} = \mathbf{d}) = 1 - \epsilon$ for some arbitrarily small $\epsilon > 0$. Then $P(X, Y, \mathbf{D} | \mathbf{C}) \approx P(X, Y | \mathbf{C}) = P(X, Y)$, and if X and Y are strongly correlated, it becomes impossible to find a compatible distribution under which $X \perp\!\!\!\perp Y | \mathbf{C}$. However, since the only dependence is between X and Y , such a distribution would certainly obey the global Markov property with respect to \mathcal{G}' , which contains the edge $X \rightarrow Y$. \square

Remark 4.4. *In other words, the Corollary states there exists some non-trivial (i.e. falsifiable) condition on the joint distribution which must be satisfied under \mathcal{G} , but not necessarily under \mathcal{G}' . In many cases we can choose smaller sets \mathbf{D} than the one used in the proof of Corollary 4.3; the generated inequalities will tend to be more powerful if \mathbf{D} is smaller, so certainly a minimal set should be used.*

It is important to stress that this result is not a causal one, and the constraints are merely a consequence of marginalizing distributions obeying certain conditional independence constraints. In the next subsection, however, we will extend this method to estimate the strength of causal relationships.

In the IV graph in Figure 1(a), Z is e-separated from Y after deletion of X , giving an inequality constraint. In general, the additional constraint implied by

the Theorem may be an inequality or a conditional independence (if $\mathbf{D} = \emptyset$); an inequality constructed will in some cases be a weaker manifestation of a Verma constraint, or possibly some other as yet unknown form of equality constraint. Verma constraints are still poorly understood; see Tian and Pearl (2002) for methods on deriving them.

Remark 4.5. *Theorem 4.2 can be extended to continuous state spaces without difficulty, but it is necessary for the set \mathbf{D} to contain only discrete variables. The IV model from Figure 1(a) with continuous X is unconstrained, for example.*

4.1 Causal Bounds

As with the IV model, we can find bounds on the average controlled direct effect due to the edge $X \rightarrow Y$ in arbitrary models, so long as X and Y are not directly confounded. First we generalize the average controlled direct effect slightly to allow conditioning:

$$\text{ACDE}_{X \rightarrow Y}(\mathbf{d} | \mathbf{c}) \equiv p(y_1 | \text{do}(x_1, \mathbf{d}), \mathbf{c}) - p(y_1 | \text{do}(x_0, \mathbf{d}), \mathbf{c}).$$

In general $\text{ACDE}_{X \rightarrow Y}(\mathbf{d} | \mathbf{c}) \neq \text{ACDE}_{X \rightarrow Y}(\mathbf{d}, \mathbf{c})$, but for appropriate graphs if $\text{ACDE}_{X \rightarrow Y}(\mathbf{d} | \mathbf{c}) \neq 0$, then $X \rightarrow Y$.

Theorem 4.6. *Let \mathcal{G} be a DAG containing the edge $X \rightarrow Y$ and observable sets of vertices \mathbf{C} , \mathbf{D} such that no vertex in \mathbf{C} is a descendant of one in \mathbf{D} . Suppose further that if the edge $X \rightarrow Y$ is removed, X is e -separated from Y by \mathbf{C} after deletion of \mathbf{D} .*

Let

$$L(x, y, \mathbf{d} | \mathbf{c}) = \max \left\{ 0, \frac{p(x, y, \mathbf{d} | \mathbf{c})}{p(x, \mathbf{d} | \mathbf{c}) + 1 - p(\mathbf{d} | \mathbf{c})} \right\}$$

$$U(x, y, \mathbf{d} | \mathbf{c}) = \min \left\{ \frac{p(x, y, \mathbf{d} | \mathbf{c}) + 1 - p(\mathbf{d} | \mathbf{c})}{p(x, \mathbf{d} | \mathbf{c}) + 1 - p(\mathbf{d} | \mathbf{c})}, 1 \right\}.$$

Then

$$L(x, y, \mathbf{d} | \mathbf{c}) \leq p(y | \text{do}(x, \mathbf{d}), \mathbf{c}) \leq U(x, y, \mathbf{d} | \mathbf{c})$$

and consequently for binary X and Y ,

$$L(x_1, y_1, \mathbf{d} | \mathbf{c}) - U(x_0, y_1, \mathbf{d} | \mathbf{c}) \leq \text{ACDE}_{X \rightarrow Y}(\mathbf{d} | \mathbf{c}) \leq U(x_1, y_1, \mathbf{d} | \mathbf{c}) - L(x_0, y_1, \mathbf{d} | \mathbf{c}).$$

If in addition X is not a descendant of any vertex in \mathbf{D} , these inequalities can be strengthened using

$$L(x, y, \mathbf{d} | \mathbf{c}) = p(y, \mathbf{d} | x, \mathbf{c})$$

$$U(x, y, \mathbf{d} | \mathbf{c}) = p(y, \mathbf{d} | x, \mathbf{c}) + 1 - p(\mathbf{d} | x, \mathbf{c}).$$




Figure 3: The unrelated confounding (UC) model.




Figure 4: (a) A DAG with three independent unobserved variables; we have avoided explicitly drawing a vertex for each of the three unobserved variables, and instead use a bidirected (\leftrightarrow) edge to indicate its two (observed) children. (b) The same graph after deletion of W .

Proof. See appendix. \square

Remark 4.7. *This result shows that we can always bound the effect corresponding to a directed edge, at least for some observed distributions, provided the two variables involved are not directly confounded with one another. The bounds for the ACDE include zero if the compatibility requirement from Theorem 4.2 is satisfied. If they exclude zero, then the edge $X \rightarrow Y$ must be present in the graph (given the other assumptions).*

5 Examples

The graph in Figure 3, which we refer to as the *unrelated confounding* (UC) model, has no edge between Z and Y , and nor are these two variables directly confounded. Theorem 4.2 and Corollary 4.3 therefore tell us that in the discrete case, the joint distribution of (X, Y, Z) is restricted, and in particular that for each ξ , the joint probabilities $p(\xi, y, z)$ must be compatible with a distribution in which $Z \perp\!\!\!\perp Y$. Let $p_{ijk} \equiv p(x_i, y_j, z_k)$; in the binary case, given $p_{000}, p_{010}, p_{001}, p_{011}$, we need to find non-negative $p_{100}^*, p_{110}^*, p_{101}^*, p_{111}^*$ such that

$$(p_{000} + p_{100}^*)(p_{011} + p_{111}^*) = (p_{010} + p_{110}^*)(p_{001} + p_{101}^*)$$

and $\sum_{jk} p_{0jk} + \sum_{jk} p_{1jk}^* = 1$. This will not be possible if, for example, p_{000} and p_{011} are both large; that is, we cannot have both $P(X = \xi)$ be large and Z and

Y strongly correlated conditional on $X = \xi$. Unlike in the IV model we cannot apply the stronger condition of Theorem 4.2, because Z is a descendant of X . We remark that (observationally) the UC model strictly contains the IV model in Figure 1(a). Note that a linear programming approach to finding constraints on this graph is not possible, so the constructive nature of the proof of Theorem 4.2 is crucial in determining how we can test this model.

The graph in Figure 4(a) is constrained in the discrete case because there is no edge between X and Y . Specifically X is e-separated from Y given W after deletion of Z , and also given Z after deletion of W (the latter being illustrated in Figure 4(b)). Note that X is a descendant of W , but not of Z , so the bounds given by Theorem 4.6 are not symmetric in the two cases. For example:

$$\begin{aligned} p(y \mid \text{do}(x, w), z) &\leq p(y, w \mid x, z) + 1 - p(w \mid x, z) \\ p(y \mid \text{do}(x, z), w) &\leq \frac{p(x, y, z \mid w) + 1 - p(z \mid w)}{p(x, z \mid w) + 1 - p(z \mid w)}. \end{aligned}$$

The first bound is likely to be stronger, though this will not hold in all cases.

6 Discussion

We have presented a graphical approach to finding inequality constraints in distributions corresponding to marginalized DAGs, based on the e-separation criterion. It can be shown that the bounds derived from the algorithm of Kang and Tian (2006) also imply the causal constraints given in Theorem 4.6, however that approach involves listing exponentially many inequalities and then using Fourier-Motzkin elimination to derive bounds. For even modestly sized graphs this becomes infeasible because Fourier-Motzkin is doubly-exponential in the number of variables in the elimination.

The advantage of the results given above is that they are ‘off the shelf’, in the sense that we need only check the conditions of the Theorems and then apply the results. Exhaustively searching possible sets \mathbf{C} and \mathbf{D} would be computationally intensive, but in many cases it is likely that good heuristics could be obtained for their selection. This could be highly advantageous in systems with large numbers of variables, especially during computationally intensive model search procedures. A further benefit of the e-separation criterion is that it is much easier and more intuitive for a human user to apply than using the algorithm of Kang and Tian (2006).

The bounds derived from Theorem 4.2 are known not to be tight in some cases, including the IV model when the instrument takes three or more states. However finding constraints from marginalized models is computationally intensive, even if the inequalities are linear, so a fast method for finding a subset of conditions may be very useful in practice.

A Proofs

Proof of Theorem 4.2. By the global Markov property for DAGs, the joint distribution P over the observed variables takes the form (1). Now, for each factor $f(V \mid \text{pa}_{\mathcal{G}}(V))$, construct a new conditional density $f^*(V \mid \text{pa}^*(V))$ where $\text{pa}^*(V) = \text{pa}_{\mathcal{G}}(V) \setminus \mathbf{D}$, by fixing any element of $\mathbf{D} \cap \text{pa}_{\mathcal{G}}(V)$ to the value specified by $\mathbf{D} = \mathbf{d}$. Note we only fix elements in the conditioning set, so f^* is still a valid conditional density.

Then the joint distribution P^* given by

$$\int_{U \in \mathcal{U}} \prod_{V \in \mathcal{V}} f^*(V \mid \text{pa}^*(V)) dU$$

factorizes according to the DAG \mathcal{G}^* formed by removing any edges in \mathcal{G} which originate in \mathbf{D} (i.e. the non-arrowhead end is incident to a vertex in \mathbf{D}). By Lemma 4.1, \mathbf{A} and \mathbf{B} are d-separated by \mathbf{C} in \mathcal{G}^* , and therefore the global Markov property for DAGs says that $\mathbf{A} \perp\!\!\!\perp \mathbf{B} \mid \mathbf{C} [P^*]$. Further, $P(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}) = P^*(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d})$ for the fixed $\mathbf{D} = \mathbf{d}$ and any $\mathbf{a}, \mathbf{b}, \mathbf{c}$, and $P^*(\mathbf{c}) = P(\mathbf{c})$ because the distribution of vertices ordered before \mathbf{D} will be unchanged. This gives the compatibility condition. If \mathbf{A} is also ordered before \mathbf{D} then $P^*(\mathbf{a}, \mathbf{c}) = P(\mathbf{a}, \mathbf{c})$, giving the stronger condition. \square

Proof of Theorem 4.6. For simplicity we will assume $\mathbf{C} = \emptyset$, but the extension to the general case is easy. Let P^* be the distribution formed by fixing $\mathbf{D} = \mathbf{d}$ in conditioning sets in the factorization of P , as in the proof of Theorem 4.2. Then

$$\begin{aligned} p(y \mid \text{do}(x, \mathbf{d})) &= p^*(y \mid x) = \frac{p^*(y, x)}{p^*(x)} \\ &= \frac{p(y, x, \mathbf{d}) + \sum_{\mathbf{d}' \neq \mathbf{d}} p^*(y, x, \mathbf{d}')}{p(x, \mathbf{d}) + \sum_{\mathbf{d}' \neq \mathbf{d}} p^*(x, \mathbf{d}')} \end{aligned}$$

Clearly $\sum_{\mathbf{d}' \neq \mathbf{d}} p^*(y, x, \mathbf{d}') \leq \sum_{\mathbf{d}' \neq \mathbf{d}} p^*(x, \mathbf{d}') \leq 1 - p(\mathbf{d})$; the expression is maximized by both these sums taking their largest possible values, and minimized when the first is zero and the second is $1 - p(\mathbf{d})$. This gives the main result. If X is not a descendant of \mathbf{D} we have $p^*(x) = p(x)$, and arrive at the tighter bounds by a similar analysis. \square

References

- A. Balke and J. Pearl. Bounds on treatment effects from studies with imperfect compliance. *JASA*, 92(439):1171–1176, 1997.
- B. Bonet. Instrumentality tests revisited. In *UAI-01*, pages 48–55, 2001.
- Z. Cai, M. Kuroki, J. Pearl, and J. Tian. Bounds on direct effects in the presence of confounded intermediate variables. *Biometrics*, 64(3):695–701, 2008.

- C. Kang and J. Tian. Inequality constraints in causal models with hidden variables. In *UAI-06*, pages 411–420, 2006.
- J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, second edition, 2009.
- J. Pearl. On the testability of causal models with latent and instrumental variables. In *UAI-95*, volume 11, pages 435–443, 1995.
- R.R. Ramsahai. Causal bounds and observable constraints for non-deterministic models. *JMLR*, 13:829–848, 2012.
- J. M. Robins, VanderWeele T. J., and T. S. Richardson. Comment on ‘Causal effects in the presence of non compliance: a latent variable interpretation’ by Antonio Forcina. *Metron*, 64(3):288–298, 2006.
- J. Tian and J. Pearl. On the testable implications of causal models with hidden variables. In *UAI-02*, pages 519–527, 2002.
- T. S. Verma and J. Pearl. Causal networks: semantics and expressiveness. In *UAI-88*, pages 352–359, 1988.
- T. S. Verma and J. Pearl. Equivalence and synthesis of causal models. In *UAI-91*, pages 255–268, 1991.